

WHAT IS CLAIMED IS:

1. A computer-readable medium having thereon computer-executable instructions for performing a method comprising determining, at a peer, to join an overlay network peer group if a first set of transport network distances is near to a second set of transport network distances, the first set of transport network distances comprising at least one transport network distance between the overlay network peer group and at least one overlay network peer group neighbor of the overlay network peer group, and the second set of transport network distances comprising at least one transport network distance between the peer and said at least one overlay network peer group neighbor of the overlay network peer group.

2. The computer-readable medium of claim 1, wherein the first set of transport network distances is near to the second set of transport network distances if, for each transport network distance in the first set of transport network distances, an absolute value of a difference between the transport network distance in the first set of transport network distances and a corresponding transport network distance in the second set of transport network distances is less than a threshold value.

3. The computer-readable medium of claim 1, wherein a transport network distance between a first node and a second node in a transport network comprises a round-trip time for a message between the first node and the second node.

4. The computer-readable medium of claim 1, wherein a transport network distance between a first node and a second node in a transport network comprises transport network latency between the first node and the second node.

5. The computer-readable medium of claim 1, wherein a transport network distance between a first node and a second node in a transport network comprises a count of transport network routing hops between the first node and the second node.

6. The computer-readable medium of claim 1, wherein:
each overlay network peer group comprises a peer group leader;

the transport network distance between the peer and an overlay network peer group comprises the transport network distance between the peer and the peer group leader of the overlay network peer group; and

the transport network distance between a first overlay network peer group and a second overlay network peer group comprises the transport network distance between the peer group leader of the first overlay network peer group and the peer group leader of the second overlay network peer group.

7. The computer-readable medium of claim 1, wherein the method further comprises:

querying the overlay network peer group for the first set of transport network distances; and

measuring each transport network distance in the second set of transport network distances.

8. The computer-readable medium of claim 1, wherein the method further comprises:

if the peer does not join the overlay network peer group, adding the at least one overlay network peer group neighbor of the overlay network peer group to a list of candidates; and selecting the nearest overlay network peer group in the list of candidates as the next to be considered for joining.

9. The computer-readable medium of claim 8, wherein the method further comprises determining to establish a new overlay network peer group if, after testing each selected candidate, the peer has not joined an existing overlay network peer group.

10. A computer-readable medium having thereon computer-executable instructions for performing a method comprising grouping overlay network peers such that each peer in a peer group has a similar transport network proximity measure with respect to peers in other peer groups.

11. The computer-readable medium of claim 10, wherein the transport network proximity measure comprises a communications round-trip time between peers.

12. The computer-readable medium of claim 10, wherein similar transport network proximity measures have an absolute difference less than a threshold value.

13. The computer-readable medium of claim 10, wherein the grouping of the overlay network peers is performed by the overlay network peers.

14. The computer-readable medium of claim 10, wherein grouping overlay network peers comprises:

adding a new overlay network peer to an existing peer group if transport network proximity measures between the existing peer group and peer group neighbors of the existing peer group are similar to transport network proximity measures between the new overlay network peer and the peer group neighbors of the existing peer group; and

if the new overlay network peer is not added to an existing peer group, establishing a new peer group comprising the new overlay network peer.

15. The computer-readable medium of claim 14, wherein establishing a new peer group comprises:

generating a new peer group identifier; and

establishing at least one overlay network connection to at least one peer group that is nearby in the transport network.

16. The computer-readable medium of claim 15, wherein:

each peer group comprises a peer group leader; and

establishing an overlay network connection between a first peer group and a second peer group comprises establishing the overlay network connection between the leader of the first peer group and the leader of the second peer group.

17. The computer-readable medium of claim 14, wherein:

each peer group has at least one neighboring peer group;

and

before establishing a new peer group, the new overlay network peer traverses existing peer groups by a method comprising:

adding the at least one neighboring peer group of the current peer group to a list of candidate peer groups; and
selecting a next peer group from the list of candidate peer groups that is at a minimum transport network distance from the new overlay network peer.

18. A computerized system, comprising a join locality-aware overlay module configured to, at least, determine that an overlay network peer should join an overlay network peer group if a first set of transport network distances is near to a second set of transport network distances, the first set of transport network distances comprising at least one transport network distance between the overlay network peer group and at least one overlay network peer group neighbor of the overlay network peer group, and the second set of transport network distances comprising at least one transport network distance between the overlay network peer and the at least one overlay network peer group neighbor of the overlay network peer group.

19. The computerized system of claim 18, wherein the first set of transport network distances is near to the second set of transport network distances if, for each transport network distance in the first set of transport network distances, an arithmetic absolute value of a difference between the transport network distance in the first set of transport network distances and a corresponding transport network distance in the second set of transport network distances is less than a threshold value.

20. The computerized system of claim 18, wherein the join locality-aware overlay module is further configured to, at least:

query the overlay network peer group for the first set of transport network distances; and

measure each transport network distance in the second set of transport network distances.

21. The computerized system of claim 18, further comprising an intra-group cache comprising an ordered leadership list listing at least one overlay network peer in an overlay network peer group that will become, in the listed order, a leader of the overlay network peer group.

22. The computerized system of claim 18, further comprising an intra-group maintenance module configured to, at least, determine if a current leader of an overlay network peer group has left the overlay network peer group.

23. The computerized system of claim 18, further comprising an inter-group cache comprising:

a list of at least one overlay network peer group neighbor of an overlay network peer group; and

for each neighbor in the list, a measured transport network distance between the overlay network peer group and the neighbor.

24. The computerized system of claim 18, further comprising an inter-group maintenance module configured to, at least, periodically measure a transport network distance

between an overlay network peer group and each overlay network peer group neighbor of the overlay network peer group.

25. A computer-readable medium having thereon computer-executable instructions for performing a method comprising:

from a first overlay network peer group, querying a second overlay network peer group for at least one overlay network peer group neighbor of the second overlay network peer group;

measuring a transport network distance between the first overlay network peer group and each of said at least one overlay network peer group neighbor of the second overlay network peer group; and

establishing at least one overlay network connection between the first overlay network peer group and one of said at least one overlay network peer group neighbor of the second overlay network peer group at a minimum measured transport network distance from the first overlay network peer group.

26. The computer-readable medium of claim 25, wherein each overlay network peer group comprises a peer group leader.

27. The computer-readable medium of claim 26, wherein querying the second overlay network peer group comprises querying the peer group leader of the second overlay network peer group.

28. The computer-readable medium of claim 26, wherein querying from the first overlay network peer group comprises querying from the peer group leader of the first overlay network peer group.

29. The computer-readable medium of claim 26, wherein measuring a transport network distance between a pair of overlay network peer groups comprises measuring the transport network distance between the peer group leaders of the pair of overlay network peer groups.

30. The computer-readable medium of claim 26, wherein establishing an overlay network connection between a pair of overlay network peer groups comprises establishing the overlay network connection between the peer group leaders of the pair of overlay network peer groups.

31. The computer-readable medium of claim 25, wherein a transport network distance between a pair of overlay network peer groups comprises a round-trip time for a message between the pair of overlay network peer groups.

32. A computer-implemented method, comprising determining to join an overlay network peer group if a first set of transport network distances is near to a second set of transport network distances, the first set of transport network distances comprising at least one transport network distance between the overlay network peer group and at least one overlay network peer group neighbor of the overlay network peer group, and the second set of transport network distances comprising at least one transport network distance between a peer and said at least one overlay network peer group neighbor of the overlay network peer group.

33. The method of claim 32, further comprising:

querying the overlay network peer group for the first set of transport network distances; and

measuring each transport network distance in the second set of transport network distances.

34. A computer-readable medium having thereon computer-executable instructions for performing the method of claim 32.

35. A computer-implemented method, comprising grouping overlay network peers such that each peer in a peer group has a similar transport network proximity measure with respect to peers in other peer groups.

36. The method of claim 35, wherein the grouping of the overlay network peers is performed by the overlay network peers.

37. A computer-readable medium having thereon computer-executable instructions for performing the method of claim 35.

38. A computer-implemented method, comprising:
querying a first overlay network peer group for one or more overlay network peer group neighbors of the first overlay network peer group;

measuring a transport network distance between a second overlay network peer group and each of the one or more overlay network peer group neighbors of the first overlay network peer group; and

establishing at least one overlay network connection between the second overlay network peer group and one of the one or more overlay network peer group neighbors of the first

overlay network peer group that is at a minimum transport network distance from the second overlay network peer group.

39. The method of claim 38, wherein each overlay network peer group comprises an overlay network peer that is a leader of the overlay network peer group.

40. The method of claim 39, wherein a transport network distance between a pair of overlay network peer groups comprises a round-trip time for a message between the overlay network peers that are the leaders of the pair of overlay network peer groups.

41. A computer-readable medium having thereon computer-executable instructions for performing the method of claim 38.